

# CIRCUIT CELLAR®

THE MAGAZINE FOR COMPUTER APPLICATIONS

## CONSIDERING THE DETAILS

Bob Perrin

## I/O For Embedded Controllers

### Part 1: Digital I/O

I/O, I/O, so off to work.... Designing generic controllers with only guesstimations of what the end-product I/O needs might be keeps Bob busy at work. When it comes to designing I/O, save yourself some work and take a look at the circuits that he covers in the first half of this series.



Designing generic embedded controllers is as much of an art as it is a science.

Every company I've worked for has attempted to spin a controller board that could be reused in future projects. These companies range from a scientific research instrumentation company to a company that built wafer-handling tools for VLSI fab. These efforts have met with varying degrees of success.

Currently, I work for a company that makes C-programmable embedded controllers. Designing commercially successful generic controllers is an interesting problem. You have to try to predict your customer's applications and deduce their I/O requirements—no meager task.

If you work for a company that wants to develop an embedded controller as a base for current and future products, you're lucky. This type of project is just plain fun.

Being intimately familiar with your company's product line will help you predict what type of applications the controller will be applied to. Knowing this, you should be able to make some

educated guesses about the I/O mix you'll need.

The I/O circuits presented here have shown themselves to be useful in a wide variety of applications. The next time you have to design I/O for your firm's embedded controller project, perhaps one or more of the circuits or ideas presented here will be of use.

In Part 1 I focus on digital techniques, and the next installment will address analog signal conditioning, A/D conversion, and D/A conversion.

### DIGITAL INPUTS

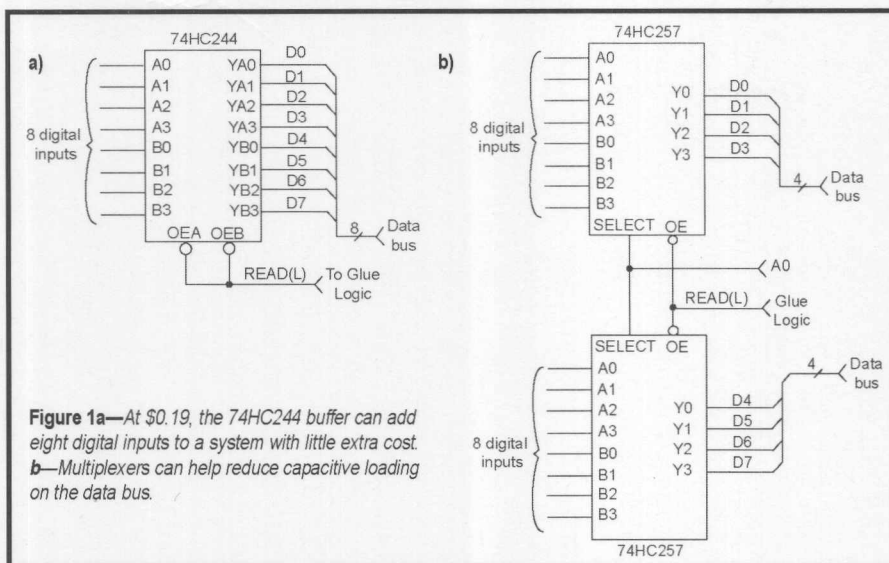
One of the simplest forms of I/O is a digital input. There are several common implementations and each has a unique set of tradeoffs.

The simplest form of digital input is shown in Figure 1a. The 74HC244 buffer sits between the processor and the outside world. When the processor wants to read the status of devices on the input port, the 74HC244's output enable is asserted, and data flows through the buffer and onto the data bus.

If the system must have many digital inputs, the 74HC244 scheme may add an unacceptable level of capacitance to the microprocessor's data bus. The tristated output of the 74HC244 has a worst-case capacitance of 20 pF ([see the 74HC244 datasheet](#)). It doesn't take many 74HC244s on a bus to slow it down.

One alternative to the 74HC244 scheme is to use a multiplexer such as the 74HC257 shown in Figure 1b. The 74HC257 has a maximum output capacitance of 15 pF ([see the 74HC257 datasheet](#)). This capacitance is slightly less than the 74HC244, and when you consider that this scheme gives two inputs per data bus line, the equivalent loading is 7.5 pF per input (compared to 20 pF for the 74HC244).

The 74HC151 is an 8-to-1 mux, but it doesn't have the ability to tristate its



outputs. However, these could be used in conjunction with 74HC244s to provide digital inputs with only one-eighth the capacitive loading per digital input that would be incurred with only 74HC-244s.

The 74HC244 or 74HC257 schemes work if you don't need to simultaneously sample more digital inputs than the data bus can carry at one time. If you do need to simultaneously sample a larger number of inputs, you'll need to latch the data with any number of latches (see Figure 1c). Some of the more popular choices are the 74HC374 and 74HC574.

The two parts are functionally equivalent, but the 74HC-574 physically has all of the inputs on one side of the IC and all of the outputs on the other side. This arrangement can make PCB routing simpler. However, at 15 pF, the output capacitance of the 74HC574 is almost as high as on the 74HC-244.

Once again, multiplexers could be used to reduce capacitive loading per input ratio. If muxes are used, the tradeoff is that you limit the number and combination of inputs that can be sampled simultaneously.

There is a more elegant solution. Figure 1d shows how 74HC597 shift registers can be cascaded to provide a large number of digital inputs with a

minimum amount of capacitive bus loading.

The 74HC597 has eight flip-flops connected to the inputs. These form an edge-triggered input latch. The 74HC-597 also has eight additional edge-triggered flip-flops that comprise a shift register.

In Figure 1d, data is simultaneously sampled on all inputs when the glue logic causes a rising edge on the RCLK signal.

Next, the processor, through the glue logic, commands the data to be moved from the input latches into the shift register. This is done by asserting

SLOAD(L).

With the data in the shift register, the processor has only to clock the data through one bit at a time using SCLK. The data is read by asserting READ(H) and looking at D0.

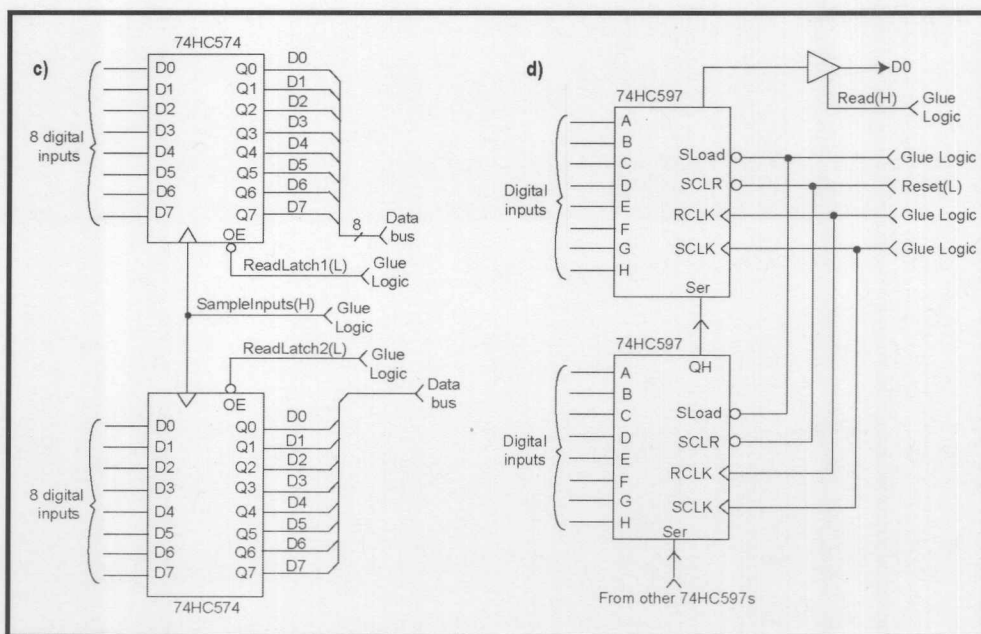
A simple variant of this circuit involves bringing the QH signals through multiple buffers, say a 74HC244, onto separate data bus lines. This process reduces the amount of time required to retrieve the serial data because data can now be read one word at a time.

Now that we have a few methods for interfacing the microprocessor's data bus to the outside world, it's time to consider the practical aspects of protecting the system from the abuse that comes with connections to the outside world.

## PROTECTING THE DIGITAL INPUTS

It's generally considered bad design practice to leave unconnected CMOS inputs floating. Unconnected inputs have a tendency to bang between input thresholds ( $V_{IL}$  and  $V_{IH}$ ). The result is that the internal transistors spend a lot of time switching unnecessarily. Not only does this contribute to noise in the system, but it consumes real power.

Installing a high-value pull-up or pull-down resistor ensures that any unconnected inputs are pulled to a known level. CMOS inputs usually



have maximum input currents of 1  $\mu\text{A}$ , with typical values much smaller. So, you can use resistors approaching 1 megaohm as pullups or pulldowns. I generally use 100 kilohms to further reduce the EMI susceptibility of the input.

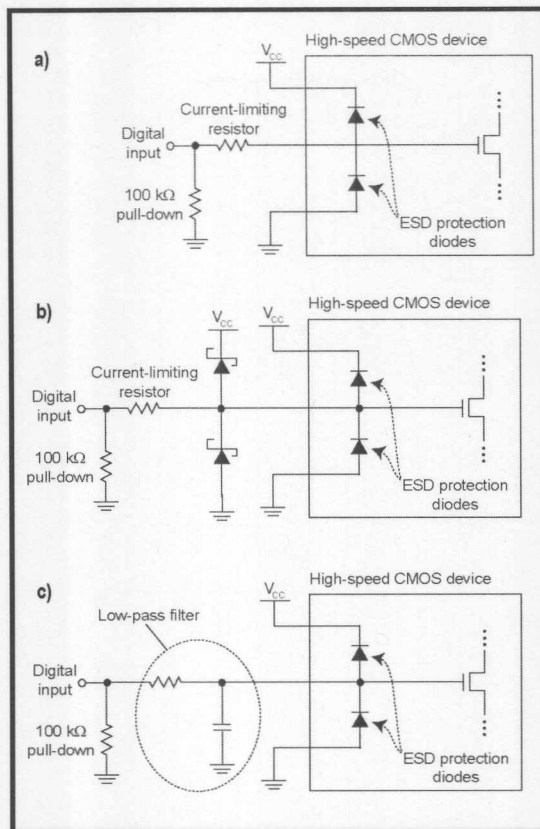
In many embedded applications, the inputs are occasionally stressed significantly above 5 V or even below ground. A simple series resistor can protect digital inputs from such overvoltages. Figure 2a shows such an arrangement. The internal diodes will clamp the voltage at the input of the CMOS device. These diodes are part of the ESD protection scheme designed into high-speed CMOS (74HCxxx) devices.

As long as the current into the input is sufficiently limited, no damage occurs to the device. The input-protection diodes are fabricated to mitigate or eliminate ESD damage to CMOS devices during the board-level manufacturing process. However, these diodes can also be used as clamps for overvoltage protection.

A more conservative and more expensive design, shown in Figure 2b, adds external discrete Schottky diodes. The forward voltage drop on these devices is about one-third of the silicon diodes in the IC's ESD protection circuit. Therefore, the internal diodes will never conduct. All the current will be carried by the forward-biased discrete Schottky diodes.

This type of external overvoltage protection uses precious PCB real estate, and besides component cost, there's also an insertion cost to populate these components. For small passive components, the insertion cost will be dominant and can't be ignored. Hewlett-Packard makes a dual Schottky diode in a SOT-23.

In Figure 2b, the external Schottky diodes prevent any possibility of CMOS latch-up on the IC. There will also be an increased immunity to ESD. The level of robustness demanded by the application dictates whether all, some, or none of the inputs require this added protection.



**Figure 2a**—A simple resistor coupled with the internal ESD protection diodes of a high-speed CMOS device makes a surprisingly robust digital input. **b**—For the cautious, Schottky diodes eliminate any chance of CMOS latch-up. **c**—A capacitor along with the current-limiting input resistor can form an RC filter to debounce the input.

Figure 2c shows a capacitor at the input of the CMOS gate. This setup serves two purposes. First, the RC circuit forms a low-pass filter. This acts to debounce the input. Glitches or contact bounces will be filtered out. Secondly, this low-pass filter also provides added immunity to ESD.

If the capacitor were ideal, a 0.1- $\mu\text{F}$  capacitor behind a 22-kilohm resistor would be adequate ESD protection for any application. Unfortunately, the finite ESR (effective series resistance) and ESL (effective series inductance) of the physical device prevent the circuit from performing optimally. Figure 2d shows how ESR and ESL are modeled.

Most capacitor manufacturers characterize typical ESR and ESL behavior of their devices, enabling you to model the behavior of the circuit. But, modeling is a bit of an art when it comes to predicting a physical circuit's behavior during an ESD event. For example, the input protection diodes in most CMOS devices are preceded by a polysilicon

resistor of unspecified value. And one of the diodes is often a distributed "diode resistor" (see the [An-313 datasheet](#)). These, in conjunction with difficult-to-predict ESD discharge waveforms, make modeling a dubious tool at best.

Unfortunately, testing is the only way I know to get a feel for how well a circuit will stand up to ESD. Testing can be risky, and lab techniques are somewhat subjective. Even testing a couple dozen devices is an incredibly small sample set to extrapolate performance data from, but I know of no better way.

Our engineering group uses a Schaffner NSG-435 ESD gun to simulate transient events. This \$7000 piece of lab equipment enables us to zap our circuits with up to  $\pm 16.5\text{-kV}$  events. We have had great success weeding out weak designs with this tool. Several companies rent this or similar ESD guns.

If your application requires that little bit of extra insurance, you can always place a transient voltage suppressor (TVS) on the input.

Many companies build TVS's, but General Semiconductor's Transorbs are probably the best known and are second-sourced by several manufacturers.

Figure 2e shows the most heavily armored input that is practical. L1 is a ferrite bead to reduce conducted RFI. L1, R1, and the open-circuit capacitance of the TVS help slow down any high-speed transient event until the TVS can turn on. R1 also serves to limit current into the TVS if the input is pulled to a low-impedance source that is of a higher potential than the TVS's standoff voltage.

R2 and C1 form a low-pass filter that will debounce switch contacts and further attenuate transient events. The diodes D1 and D2 prevent the CMOS input from going above  $V_{CC}$  or below ground by 0.2 V. D1 and D2 will also provide additional ESD protection.

The circuit shown in Figure 2e is gross overkill for most applications. However, for certain industrial or mission-critical applications, this is the



best protection you can have without resorting to optoisolation. Insertion costs, component cost, and board space are the factors to be weighed against ESD protection, overvoltage protection, and debouncing needs.

A bit of caution should be used when selecting resistors for use in circuits that may be subjected to ESD. It turns out that plain old carbon-composition axial-leaded resistors are the best [1]. Metal film axial resistors and surface-mount resistors have patterns cut into the film to trim the film geometry to achieve the desired resistance. ESD has a tendency to jump the insulative gaps etched into the metal film.

This behavior has two ramifications. First, the resistor's value is effectively reduced during the event. Second, ionization paths may form, creating a permanently altered resistance value.

Surface-mount resistors have similar problems, but they also have problems with hot spots forming in the metal film when subjected to ESD. These hot spots are caused by nonuniform current densities flowing in the metal film. The net result is that the resistor can be permanently damaged by ESD events.

There are other enhancements that can be added to digital inputs. Analog comparators can be used if precisely controlled switching thresholds are required. Comparators can be designed with or without hysteresis. And, if a good-size voltage divider is placed on

the input, thresholds way above  $V_{CC}$  or way below ground can be managed.

Optoisolators can also be used with digital inputs. These devices can be used to get several thousand volts of galvanic isolation. However, the input device must supply about 1000× the input current to an optoisolator than is required by a CMOS gate. Optoisolators—at least inexpensive ones—can be quite slow. And, you still need to protect the LED in the opto-isolator from ESD damage. Depending on the application, optoisolators may be a benefit or a detriment.

## DIGITAL OUTPUTS

For digital outputs, the bus loading tradeoffs are similar to those encountered with digital inputs. There are four

high-speed CMOS devices I want to discuss—the 74HC574, the 74HC273, the 74HC259, and the 74HC594.

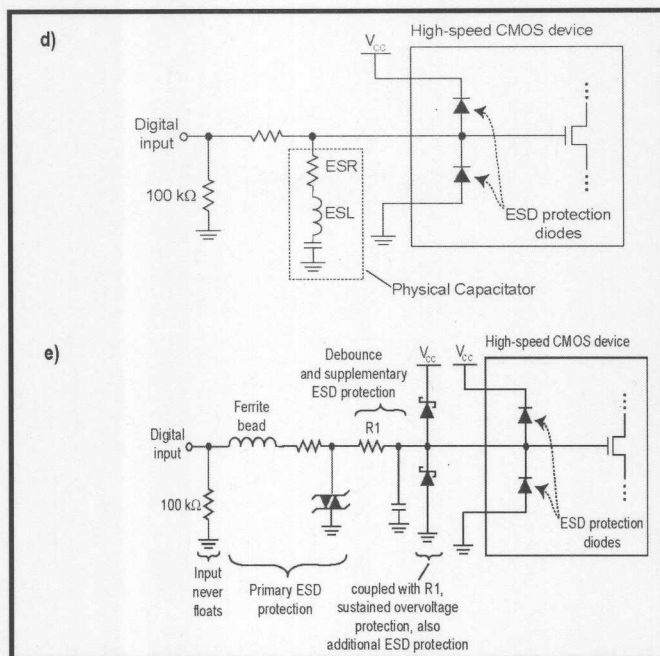
Figure 3a shows how digital outputs can be implemented with a 74HC574. This is the same part that was used to implement digital inputs in Figure 1c. Besides a high capacitive bus load per output (10 pF per output), the 74HC574 has no global reset for the latches, which means the system RESET signal can't be used to put the outputs into a known state. Having digital outputs that are nondeterministic on powerup is extremely troubling.

The circuit shown in Figure 3a works around this problem by adding an additional flip-flop and pull-down resistors on the outputs.

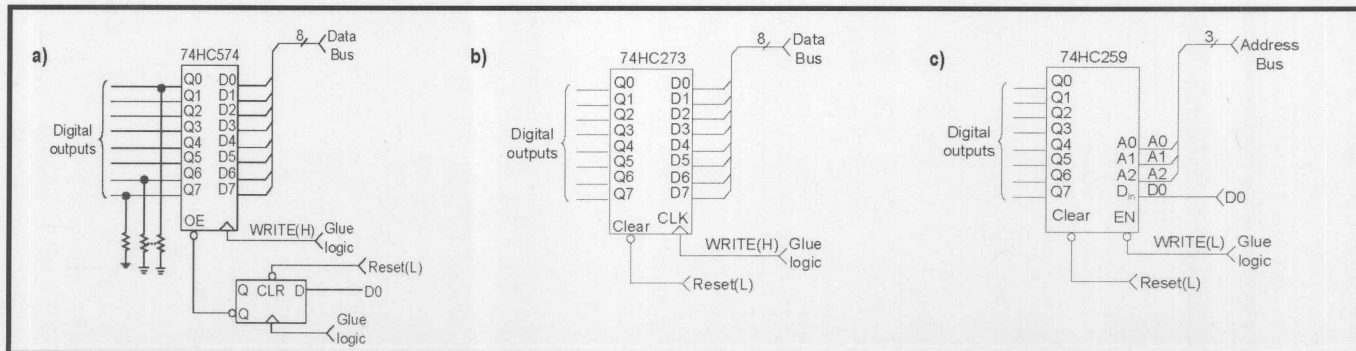
On powerup, the 74HC574's outputs are tristated and the outputs are pulled low with the 100-kilohm resistors. Once the system is up and running, the processor can write data into the 74HC574 and enable the outputs by writing a 1 into the supplementary flip-flop.

The 74HC273 has a CLR(L) signal but no output enable. This isn't a problem if your application doesn't require tristated outputs. Figure 3b shows a 74HC273 configured as an output latch. The input capacitance is 10 pF (maximum) and may present a bus loading issue if many 74HC273s are needed.

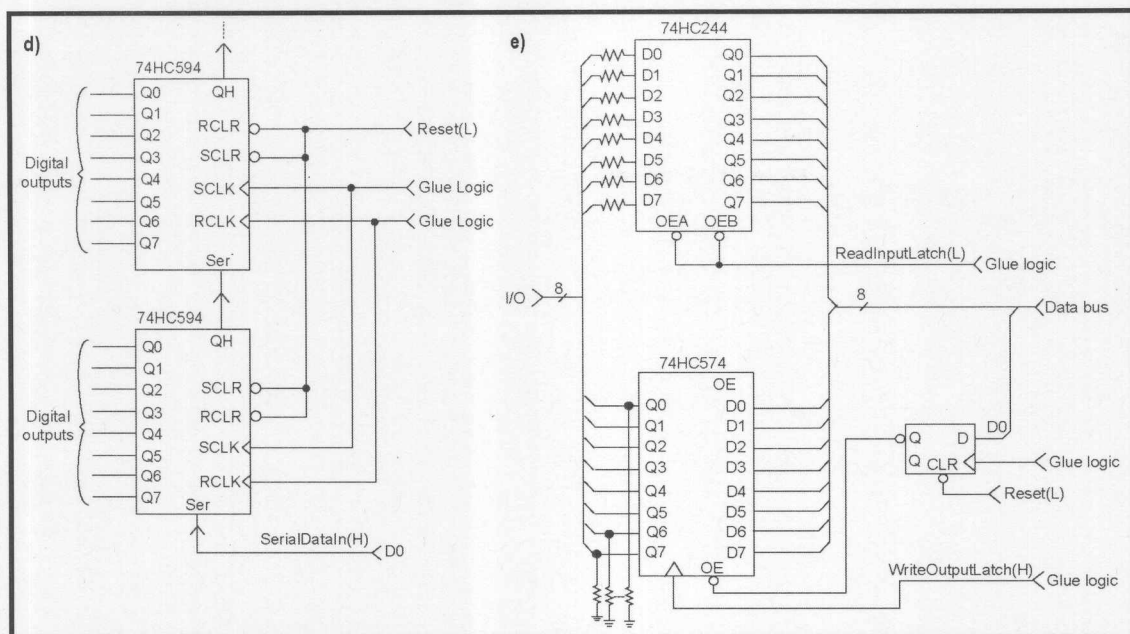
The 74HC259 is a bit-addressable latch. Figure 3c shows how to use the 74HC259 to reduce the capacitance



**Figure 2d**—ESR and ESL are unfortunate, parasitic impedances associated with physical devices. **e**—This circuit is grossly expensive on a per I/O basis, but it's also extremely rugged.



**Figure 3a**—Even at \$0.20, the lack of a CLR signal makes the 74HC574 a questionable choice for an output latch. **b**—The 74HC273 has a CLR(L) signal allowing deterministic powerup. **c**—The 74HC259 can reduce loading on the system data bus, but the microprocessor's address bus must still drive a heavy capacitive load.



**Figure 3d**—The 74HC594 is a flexible and inexpensive (\$0.25) shift register, suitable for creating digital outputs. **e**—You can combine inputs and outputs to form a set of byte-wide programmable I/O points.

load per output to 1.25 pF per output on the data bus. The 74HC259 has a requirement for three address lines, so the address bus will continue to be heavily loaded if many latches are placed on the same address lines.

One 74HC259 gotcha to look out for is the active-low level-triggered latch enable. When using the 74HC259 as an output latch, the address and data must remain valid until the rising edge (deassertion) of the WRITE(L) signal (see Figure 3c). All the other latches discussed have rising-edge-triggered clocks. The 74HC259 has a level-triggered latch enable to allow the device to be used as a 3-8 decoder in other applications.

Figure 3d shows how to use the 74HC594 shift register, which is a useful part. The device has a shift register coupled to independently controlled output latches. The data is clocked through the shift register chain and then all outputs are simultaneously updated. The 74HC594 has a clear for both the output latch and the shift register chain.

The 74HC594 does not have the ability to tristate its outputs. The 74HC595 trades the 74HC594's output latch clear pin for an output tristate pin. Other than that single pin, the 74HC594 and 74HC595 have identical pinouts.

In Figure 3d, data from D0 is clocked into the shift register chain

using SCLK. Once the shift-register chain is fully loaded, all of the outputs are updated simultaneously when the processor causes a rising edge on RCLK.

If the shift register chain is long, it can be split up into smaller chains, each fed with a separate data bus bit (D0, D1, D2, ...).

The 74HC574 has the ability to tristate the outputs, which can be useful if you need to create an I/O point that is both an input and an output. One thing to watch out for in this type of circuit is powerup initialization. If an I/O may be used as an input, it should default to an input, lest the output latch contend with an offboard digital output device. Figure 3e shows how to construct a set of I/O points.

## BRIDGING THE GAP

Now that we have some ideas on how to implement the logic associated with a digital output, let's consider how to bridge the gap between CMOS outputs and the real world. Most practical applications require more current drive capability than a CMOS digital output can deliver.


Perhaps the most obvious method is the use of a relay. Most relays require more current than a CMOS output can source (or sink). An intermediate NPN BJT is a common solution. Figure 4 shows how to drive a relay. A 1N914 is a time-honored diode for flyback sup-

pression.

Relays are usually considered bulky, but at least one company has inexpensive surface-mount devices. The Aromat TQ-SMD series parts are fantastic. They are available in many pole and throw combinations as well as in a variety of coil voltages. Latching devices are also available.

Some of the advantages mechanical relays have over other alternatives include, low contact resistance, ability to drive AC or DC loads, excellent electrical isolation, and high impedance when contacts are open. For many applications, mechanical relays are still the best fit.

Also, solid-state relays can be purchased in many forms. Solid-state relays generally carry a hefty price tag, may be fairly limited in the type of load they can drive, and may require heat sinking. However, when properly used, solid-state relays have a much longer



**Figure 4**—Relays are expensive and mechanically bulky, but for versatile high-current outputs, relays are tough to beat.

life than their mechanical cousins.

The array of devices that can be used to bridge the gap is simply too vast to cover in one article. Fortunately, if you can just make the step from the CMOS output to an output that can drive between 100 and 1000 mA, the next step—driving lots of amps—is often a just matter of selecting an appropriate contactor (relay).

Making that first step between CMOS and moderate current drive outputs can be accomplished with several devices, each with their own strengths and weaknesses.

For many years, sourcing and sinking drivers have been available in IC form. These are found in many applications, from driving LED arrays to small motors, and they are available from many manufacturers.

Motorola and Allegro sell the ULN2803, and Allegro also sells the UDN2985. The ULN2803 is a sinking driver, and the UDN2985 is a sourcing driver. Both of these devices use a Darlington pair as the output switch. And both devices have integral flyback suppression diodes on each output. Figure 5 shows the pinout and output stage for each of the two parts.

These devices have a nearly identical pinout and circuit boards can be designed to accept the parts interchangeably. To accomplish this, the PCB will also require a couple of jumpers to allow the two noninterchangeable pins (GND and K) to be jumpered to the appropriate place.

The ULN2803 can sink 500 mA per pin but is limited by the total power-dissipation capability of the package. This really means the entire package can sink around 500 mA split up across all the drivers. The maximum voltage allowed on the outputs is a respectable 50 V.

The UDN2985 can source around 250 mA. The maximum voltage drop

on the output is 30 V. And again, the maximum drive capacity is limited more by the package power-dissipation characteristics than by the Darlington transistor outputs.

When using these types of devices, one potential problem is the relatively poor ability of the Darlington pair to pull the “on voltage” near the rail. Because the output transistor in a Darlington is never driven hard into saturation, the output is only pulled within 1.2–2.5 V of the rail, depending on  $I_C$ . This can be a problem in some applications.

Let’s say you want to use a ULN2803 to drive a relay with a 5-V coil and you only have a 5-V supply. The Darlington will only allow about 3.5 V to be developed across the relay. Many 5-V relays have a maximum pick-up voltage (the voltage at which the relay is guaranteed to operate) higher than 3.5 V.

For example, the Aromat TQ-SMD relays have a maximum pick-up voltage of 3.75 V for the 5-V relay. To ensure reliable operation with a ULN2803, you have to go to the 4.5-V TQ-SMD, which has a maximum pick-up voltage of 3.38 V. The 4.5-V parts have signifi-

cantly longer lead times. Alternately, you could design in a second higher voltage power supply, say 8 V, to drive the relays, but this may add cost to the system.

The second common problem that people run into with Darlington output drivers occurs when they try to use them to drive a CMOS input. The high-voltage drop across the Darlington gives away the entire CMOS noise margin and then some.

People usually run into this kind of situation when they buy a PLC or turnkey controller that uses Darlington pairs to implement “high-current digital outputs.” When the customer tries to use these “digital outputs” to drive another device’s CMOS input, it

doesn’t work.

The real moral of the story here is, if you buy a controller, read the specifications carefully. Darlington outputs are widely used in industry and although they are versatile, they do have limitations. Ultimately, the reliability of your system rests on your shoulders, not those of your suppliers. Always dig into the specifications and schematics of off-the-shelf controllers. Many companies provide complete schematics and specifications in their manuals.

If you can’t live with the voltage drop in a Darlington, you can always just go with a BJT or MOSFET implementation of a high-current driver. Each of these has tradeoffs. First, let’s consider the BJT.

Figure 6a shows a typical NPN sinking driver circuit. The circuit is simple, but there are still a few details to keep in mind before laying this topology down in copper.

Advantages of this circuit include low cost, high  $V_{CE}$ , widely available parts, and simplicity of design. When the transistor is fully saturated,  $V_{CE(SAT)}$  can be on the order of 100–300 mV, allowing reasonably high  $I_C$ , even with physically small devices.




Figure 5—The ULN2803 and its brother, the UDN2985, are flexible parts.



The trick to making all of this work is base current  $I_B$ . The current transfer ratio (beta or  $h_{FE}$ ) in saturation is very low, on the order of 10–50. For inexpensive devices (like the MPSA-2222A or 2N3904) over temperature 10 is a safe value to use for design. There are transistors that are optimized for switching—for example, the Zetex FMMT625 SuperSOT transistors. I use a value of 20 for  $h_{FE(SAT)}$  when designing with these parts.

If you want to drive loads on the order of 500 mA with a run-of-the-mill NPN, you'll need to supply 50 mA of base current, or about 25 mA of base current if you use a more expensive (by a factor of 10) device.


Another design consideration is the total amount of current the digital latch that is used to drive the NPNs can source. A 74HC574 has an absolute maximum  $I_{CC}$  of 70 mA (Fairchild Semiconductor). The 74HC574A from Motorola has a maximum  $I_{CC}$  of 75 mA. This is the maximum current that can be pulled through the  $V_{CC}$  pin or sunk into the GND pin.

So, if you have a 74HC574 driving eight NPNs, the base current to each NPN must be limited to  $70/8 = 8.75$  mA. And that number leaves no safety margin.

The 74HC574 falls into the class of parts considered by manufacturers to need a little beefier-than-average current-handling capacity. The 74HC259, for example, is more typical of the high-speed CMOS devices and has a maximum  $I_{CC}$  of 50 mA. This translates to  $50/8 = 6.25$  mA of base current.

With a  $\beta_{(SAT)}$  of 10, a meager 6 mA  $I_B$  limits  $I_C$  to an anemic 60 mA. Unless you can supply enough base current to keep the transistor turned on hard, NPN high-current drivers don't work all that well, which is why devices like the ULN2803 use a Darlington.

Figure 6b shows how to build a PNP high-side switch. Base current in the PNP is the biggest concern. PNP high-side switches generate a lot of heat in



**Figure 6a—**This is a typical implementation of eight NPN sinking drivers built from a 74HC574. The 74HC574 can be overtaxed if all the BJTs are turned on. **b—**When building high-side drivers like this, watch out for high power dissipation in the PNP's base resistor.

their base circuits when high voltages are being switched.

If the NPN can sink 60 mA, then we should be able to have PNP  $I_C$ 's of 600 mA. That's pretty good. The trouble is, the power dissipated in the PNP's base resistor when the +RAIL voltage is high. The only way to handle this is to use a physically large resistor.

The voltage drop across the resistor is  $+RAIL - V_{BE(PNP)} - V_{CE(SAT)(NPN)}$ . For example, if the +RAIL voltage is 30 V, then the potential across the base resistor will be about 29 V (i.e.,  $30 - 0.7 - 0.3$ ).

If the current through the resistor is designed to be 60 mA, the power dissipated in the resistor will be 1.74 W. So, a 2-W resistor will be required. If you have eight of these circuits in your box, you have ( $8 \times 1.74 =$ ) 14 W of power being burned off as heat.

This condition places additional burdens on the overall system design. If you have to remove 14 W of power from your product, you may need additional ventilation or perhaps a fan.

Another issue to be aware of is the leakage current (especially at high temperatures) in the NPN in the base circuit. The pull-up resistor should be selected to be as small as possible to overcome the leakage current and keep the PNP off.

Enhancement-mode MOSFETs are a natural choice for overcoming the base drive problems associated with the BJT

circuits shown in Figure 6.

Figure 7a shows a simple low-side drive. International Rectifier and Siliconix both have nice selections of MOSFETs. Recently, I've been tinkering with the IRFL014 and IRL014N from International Rectifier. These parts are in a SOT-223 surface-mount package and cost between \$0.30 and \$0.50.

Driving these devices is much like driving a capacitor. Once the gate is charged, the leakage current from the gate to the channel is negligible in this application. The

biggest difficulty is getting the MOSFET to turn on hard with only a 5-V  $V_{GS}$ .

The less expensive IRFL014 with a 5-V gate drive will easily handle channel currents of 500 mA. The IRL014N is designed to be turned on with a 5-V  $V_{GS}$ , although the maximum allowable  $V_{GS}$  is  $\pm 15$  V.

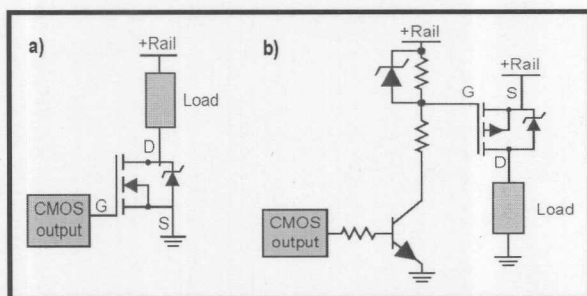
The primary disadvantage of MOSFETs over BJTs is cost. The 2N2222s are just a few pennies. SuperSOTs (like the FMMT625) are just 25 to 30 cents. MOSFETs start at 30 cents and go up rapidly.

Figure 7b shows how to use a P-channel MOSFET to build a high side switch. When +RAIL is at a relatively low voltage, we have the problem of not being able to fully turn on the MOSFET. This happens because we can't develop a high  $V_{GS}$ . At high voltages on +RAIL, we need to limit  $V_{GS}$  to avoid exceeding the maximum allowable  $V_{GS}$ .

International Rectifier's IRFL9014 is a P-channel MOSFET suitable for building low-cost, high-current drivers. At +RAIL voltages of 7–10 V, the MOSFET is turned on. The part really shines at  $V_{GS}$ 's of 15 V. The minimum  $R_{DS(ON)}$  for this device is 0.50 ohms, but the device only costs around \$0.40.

The maximum  $V_{GS}$  permitted on the IRFL9014 is 20 V. The zener shown in Figure 7b must be selected to limit  $V_{GS}$  from exceeding this maximum.

Sourcing drivers built like those



**Figure 7a**—MOSFETs require minimal gate drive current, a saving grace if you want to drive eight of them from a CMOS IC. **b**—The zener is required to prevent the maximum  $V_{GS}$  specification from being exceeded.

shown in Figure 7b can source higher currents when +RAIL is at higher voltages because we can develop a near maximum  $V_{GS}$  on the MOSFET, thus turning it on fully.

One advantage FETs have over BJTs is the ability to current share. BJTs don't current share well. You may parallel FET-based channels to your heart's content. This can be a big advantage in some applications.

If you design H-bridges for driving motors, thermoelectric devices, or other high-current devices, the above principles directly apply. Other refinements that should be considered for high-current drivers include flyback suppression for inductive loads, and ESD protection.

There are other devices available to bridge the gap between CMOS devices and the outside world. Triacs can be used if you need to switch AC loads.

## WAIT ONE, OVER

If you're designing an embedded controller for a specific application, selecting an I/O methodology is fairly straightforward. If you're designing a controller as a generic platform that's intended to meet an array of future needs, selecting an I/O mix is a bit more challenging. Lastly, if you're purchasing an off-the-shelf controller, to ensure the device will work reliably in your application, you must know how the I/O is implemented.

Buy it or build it, the reliability of the final system is up to you. The project engineer must do the analysis on the I/O systems to determine their suitability to the application at hand.

Buying turnkey controllers can reduce hardware design cycles and manufacturing overhead. However, selecting reliable interfaces suitable to the application still requires you to analyze the vendor's implementation of the I/O.

Next month, we'll explore the analog side of controller I/O.

*Bob Perrin spends his days designing general-purpose C-programmable embedded controllers and troubleshooting customer system-level problems for Z-World ([www.zworld.com](http://www.zworld.com)). Over the last ten years, Bob has designed instrumentation for agronomy, soil physics, and water activity research. He was also the lead design engineer for an intrinsically safe line of workstations for use in explosive gas and particulate environments (class 1, divisions 1 and 2). For more articles by Bob, visit his online library at [www.engineerbob.com](http://www.engineerbob.com).*

## REFERENCE

- [1] R.A. Pease, *Troubleshooting Analog Circuits*, Butterworth-Heinemann, Stoneham, MA, p. 28, 1991.

## SOURCES

### UDN2985

Allegro Microsystems, Inc.  
(508) 853-5000  
Fax: (508) 853-7861  
[www.allegromicro.com](http://www.allegromicro.com)

### TQ-SMD

Aromat Corp.  
(800) AROMAT9  
[www.aromat.com](http://www.aromat.com)

### High-speed CMOS Logic

Fairchild Semiconductor Corp.  
(800) 341-0392  
(207) 775-8100  
Fax: (207) 761-6020  
[www.fairchildsemi.com](http://www.fairchildsemi.com)

### NSF-435 ESD gun

Schaffner EMC, Inc.  
(800) 367-5566

(973) 379-7778  
Fax: (973) 379-1151  
[www.schaffer.com](http://www.schaffer.com)

### FMMT625

Zetex, Inc.  
(516) 543-7100  
Fax: (516) 864-7630

### 74HC259

Motorola, Inc.  
(561) 739-3880  
Fax: (561) 739-3815  
[www.mot.com](http://www.mot.com)

### MOSFETs

Vishay Intertechnology, Inc.  
(619) 535-9080  
Fax: (619) 535-9115  
[www.vishay.com](http://www.vishay.com)

### International Rectifier

(310) 322-3331  
Fax: (310) 322-3332  
[www.irf.com](http://www.irf.com)

## COMMENTS

Join the discussion on <http://bbs.circuitcellar.com/cgi-bin/dnewsweb.exe> and tell us what you think!

**Circuit Cellar, the Magazine for Computer Applications.** Reprinted by permission. For subscription information, call (860) 876-2199 or [subscribe@circuitcellar.com](mailto:subscribe@circuitcellar.com).



# CIRCUIT CELLAR®

THE MAGAZINE FOR COMPUTER APPLICATIONS


## CONSIDERING THE DETAILS

Bob Perrin

## I/O for Embedded Controllers

### Part 2: Analog I/O

Embedded systems are used in so many different applications, covering all possible analog I/O requirements is impossible. In Part 2 of this series on I/O, Bob offers a few circuits and components that have proven adequate for many applications in the past.



In Part 1, I covered digital inputs, outputs, and I/O points, as well as bridging the gap between relatively fragile high-speed CMOS devices and the real world. In this article, I discuss the basics of analog interfacing.

Embedded systems are used in so many different applications, it's impossible to cover all the analog I/O requirements. Here, I simply offer a few circuits and components that have proven adequate for many applications in the past.

#### ADCs

Acquiring analog signals requires an analog to digital converter (ADC). There are more ADCs on the market than can be counted in a week. When I'm in the market for an ADC, I usually look at Analog Devices, Burr-Brown, Linear Technologies, TI, and National Semiconductor, though not necessarily in that order.

I currently have a few favorite devices. For inexpensive 12-bit conversion, I'm partial to the Texas Instruments TLC2543. This device has an internal

11-channel analog multiplexer, comes in several 20-pin packages, and costs around \$4.80.

For applications requiring a fast 12-bit converter in a tiny package, I'm fond of the Analog Devices AD7887. This device is available in an 8-pin Micro-SOIC or Narrow-SOIC package. The Micro-SOIC is particularly impressive with its 0.025" pitch leads and 0.196" × 0.244" footprint.

The AD7887 has one or two analog inputs depending on how you configure the device. It will sample at 125 kilosamples per second and operate on supply rails between 2.7 and 5.25 V.

The one shortcoming of the AD7887 is the lack of an internal voltage reference. Analog Devices designed this part to be low power enough that the AD7887's  $V_{DD}$  pin can be connected directly to the output of a reference, which means the external reference can be used not only as a voltage reference but also as a voltage regulator. This scheme provides excellent decoupling from noise that may be on the raw power supply rails.

The ChipCenter Supercatalog gives a quote from Avnet of \$4.90 for the AD7887AR, which seems a bit expensive for a single-channel 12-bit converter. But given the other features the AD7887 offers, five bucks is a great deal.

There is currently a whole slew of 24-bit delta-sigma converters on the market. The AD77xx family is possibly the most well-known family of 24-bit converters. The family has parts that are intended to interface directly to several common classes of sensors.

Burr-Brown has the 24-bit ADS1211, which is a full-featured yet inexpensive part. I used this part in a project and was quite happy with the price/performance ratio. The project only required 16 noise-free bits, which was easily achieved with the ADS1211.

Burr-Brown tried to roll so many

features into the '1211 that the serial communications protocol is a bit overwhelming at first glance. There are several options for configuring data I/O lines and the self-clocked versus externally clocked protocol.

The '1211 also has many options for configuring the converter's internal sampling, filtering, calibration, and gain characteristics. All of these are controlled through the serial port. The '1211 has an internal 4-1 analog mux and comes in a 24-pin package. The mux actually has eight inputs and two outputs. Internally the ADS1211 has a true differential input, with the mux, this gives the user four channels of true differential analog input.


Although it's a sophisticated device, the '1211 is low in cost. To assist developers, Burr-Brown has published several meaty application notes focusing on various aspects of the ADS1211. There are other members in the ADS121x family. The app notes cover all variants equally well.

If physical space is a concern, I recommend looking at the LTC2400. This is a micro-power (200  $\mu$ A) delta-sigma 24-bit ADC in an 8-pin SOIC. The LTC2400 doesn't require an external oscillator, although one may be used. This ADC has one single-ended analog input. A bigger brother, the LTC2408, has eight inputs and is in a 28-pin SSOP.

The LTC2400 claims to be insensitive to PCB layout. The cute little demo board available for the LTC2400 is a two-layer board with only a single ground plane. As a further demonstration of the LTC2400's ability to deliver high-performance in a relatively noisy environment, the demo board derives its power from the RS-232 connection with its host PC. All things considered, the LTC2400 does an excellent job, in a small space, at a low cost.

## SINGLE CONDITIONING

The world is full of interesting and unique analog interfaces. In many embedded applications, the sensor data is environmental or biophysical. Both data types share the fact that they are typically slow-changing phenomena. I've found that most of the time, a simple circuit with an op-amp and a few passives will serve nicely to bridge the



**Figure 1**—This simple single-supply op-amp circuit will suffice in most voltage-scaling applications.

gap between sensor and ADC.

Figure 1 shows a versatile circuit that maps a unipolar or bipolar input voltage to a unipolar output voltage suitable for feeding an ADC. The circuit only requires a single-sided power supply.

The equation that describes the DC behavior of the circuit shown in Figure 1 is:

$$V_{OUT} = (V_{OFF} - V_{IN}) \frac{R_G}{R_{IN}} + V_{OFF}$$

where  $V_{OFF}$  is given by:

$$V_{OFF} = \left( \frac{R_{BOTTOM}}{R_{TOP} + R_{BOTTOM}} \right) V_{REF}$$

The circuit is fundamentally an inverting amplifier. The gain is given by:

$$GAIN = G = - \frac{R_G}{R_{IN}}$$

When designing the circuit, the first component to select is  $R_{IN}$ . This should be selected to be large enough not to load the sensor being measured, but small enough to allow sufficient gain.

The next component value to select is  $R_G$ . To do this, first determine the magnitude of the gain you need:

$$|G| = \frac{V_{REF}}{V_{IN_{max}} - V_{IN_{min}}}$$

This is the right equation, assuming you want to map an input voltage range into a zero to  $V_{REF}$  output range for an ADC. If this isn't the case, just substitute the magnitude of your output range for  $V_{REF}$ .

Once the gain magnitude and  $R_{IN}$  are known, use:

$$|G| = \left| - \frac{R_G}{R_{IN}} \right| = \frac{R_G}{R_{IN}}$$

Once  $R_{IN}$  and  $R_G$  are fixed, you can solve equation 1 for  $V_{OFF}$ . To do this, note that the circuit is an inverting amplifier, which means that  $V_{OUT}$  is zero when  $V_{IN}$  is at a maximum. This leads to:

$$0 = (V_{OFF} - V_{IN_{max}}) \frac{R_G}{R_{IN}} + V_{OFF}$$

$$0 = (V_{OFF} - V_{IN_{max}}) |G| + V_{OFF}$$

$$V_{OFF} = V_{IN_{max}} \times \frac{|G|}{|G| + 1}$$

Once you know  $V_{OFF}$ , you can select values for  $R_{TOP}$  and  $R_{BOTTOM}$ . I usually select 10k for  $R_{TOP}$ , then solve equation 2 for  $R_{BOTTOM}$ .

$C_1$  can be selected to add a pole at:

$$F_0 = \frac{1}{2\pi C_1 R_G}$$


Once you have standard values selected, you must verify that the standard resistor values allow the full input voltage range to map to the proper output voltage range. You can find a table of standard resistor values at [www.engineerbob.com/stdres.pdf](http://www.engineerbob.com/stdres.pdf).

Also, you must be sure that the op-amp stays linear over the output voltage range. Some op-amps have difficulty driving near the rails.

One of the biggest drawbacks of the circuit shown in Figure 1 is the fact that it's a single-ended circuit. Many sensors have a differential output. The circuit shown in Figure 2 shows a simple method to convert and scale a differential signal from a sensor to a single-ended signal suitable to hook to an ADC.

Several companies make single-chip (monolithic) instrumentation amplifiers (in-amps). Burr-Brown, Analog Devices, and Linear Technologies all have offerings. Burr-Brown's INA118 is a good place to start and Analog Devices' AD623 is a low-cost part with great features.

Monolithic in-amps often come in 8-



**Figure 2**—For applications that demand the measurement of a differential signal, the monolithic instrumentation amplifier should be considered.

pin packages and have a pinout like that shown in Figure 2. A single resistor,  $R_G$ , sets the in-amp's gain. Power, ground, two inputs, an output, and an output reference pin account for the remaining six pins.

In-amps have high-input impedances. Ten-gigaohm typical input impedances are common, which means that even high-resistance bridges can be measured with minimal loading.

The output reference pin allows the output voltage to be referenced to something other than ground. This ability can be useful in mapping a bipolar sensor's output into an ADC's input range. In Figure 2, the REF pin is connected to +REF/2. This arrangement enables the in-amp to map the bridge's bipolar output signals into a unipolar 0 to REF range.

Negative inputs are mapped between 0 and +REF/2. Positive inputs are mapped between +REF/2 and REF. This assumes that  $R_G$  is selected to appropriately scale the input signals. This arrangement is useful when the ADC's input range is 0 to +REF.

Some sensors have an output current that is proportional to measured phenomena. Unless the ADC you selected has a current input, the current must be converted to a voltage. A current-to-voltage converter is, at its simplest, a resistor.


If gain is needed, a transimpedance amplifier such as the one shown in Figure 3 may be used. The advantage of a circuit like Figure 3 is that gain may be set with a single resistor.

Several disadvantages exist as well. For example, with high gains, the circuit will exhibit a high offset. There are several variations on the basic topology shown in Figure 3 that will reduce offset and noise, at higher gains [1–3].

## DACs

Like ADCs, DACs are abundant. You can get DACs with current or voltage outputs, internal or external references, and parallel or serial digital interfaces.

DACs designed for audio applications often do not have the stability and low offset required for control systems. These devices are generally available in 16-bit resolutions at rock-bottom prices. If your application can tolerate, null, or



**Figure 3**—If you need to do an I-V conversion, a simple transimpedance amplifier is a good place to start.

otherwise compensate for these offsets, an audio DAC may be something to consider. The saving grace of audio DACs is the fact that they exhibit monotonic performance.

Cost, space, and power consumption are often the primary considerations for selecting a DAC. Analog Devices has the AD53xx family of DACs. These devices come in SOT-23 six-pin packages and can be had for a song. These parts typically consume less than 150  $\mu$ A, and operate over a 2.7–5.5-V supply range.

Table 1 compares the three devices that I have used in past projects. The data in Table 1 was extracted from the Analog Devices web site. The full AD53xx family currently has nine parts with an additional 15 parts planned. The full table can be found at [www.analog.com/support/standard\\_linear/selection\\_guides/AD53xx.html](http://www.analog.com/support/standard_linear/selection_guides/AD53xx.html)

Like any device, the AD5300 parts have pros and cons. The biggest con of these devices is initial offset. The second biggest con is the “relative accuracy” much like integral nonlinearity (INL) for ADCs. In the pros column are size, cost, scalability, 10- $\mu$ s settling time, and a 1-V/ $\mu$ s slew rate. Overall, this family of converters offers a solution for many common applications that require a digitally controllable voltage.

Signal conditioning for analog outputs is a similar problem to single conditioning for analog inputs. The circuit shown in Figure 1 is excellent for mapping a DAC output voltage to a channel output voltage.

Often an analog output is required to deliver more than the few milliamps a

jellybean op-amp can deliver. When this happens, you have several options. You can use a discrete BJT or MOSFET follower in the output stage of your circuit. This option works fine if you only need to deliver a positive or a negative voltage. But if the output is to be bipolar, a simple transistor won't work.

Class B amplifiers made from discrete parts have certain difficulties. The plethora of components is expensive. The biggest problem I've had is crossover distortion. The easiest way around all this discrete design is to buy an op-amp that has the drive capacity you need.

The Burr-Brown OPA548 is great for creating high-current analog outputs. The device comes in a 7-pin TO-220 or a 7-pin D<sup>2</sup>PAK. The part can deliver continuous 3 A of current and a peak of 5 A. Burr-Brown's “budgetary 1000-piece pricing” is \$5.45. One drawback of the OPA548 is that it requires a minimum supply rail of 8 V (or  $\pm 4$  V in a split-rail system).

The OPA548 gives the designer the ability to limit the maximum current with a single resistor. The current limit can also be digitally programmed using an external DAC.

Another important detail to consider when designing analog output stages is power-up state. For example, say you have a DAC that delivers 0–5 V feeding a circuit (similar to that shown in Figure 1) that maps the 0–5 V to  $\pm 10$  V. The power-up state of the DAC may be 0 V. That will map to a +10-V channel output.

If that isn't acceptable, one solution is to add a tristate to the output with a relay. Figure 4 shows how this scheme works. The pull-down resistor holds the channel's output at ground when the relay is open.

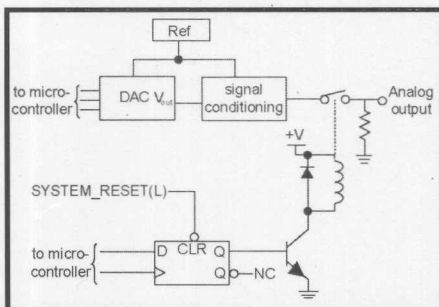
Another solution is to use a flip-flop and a small MOSFET connected to the system RESET to pull a critical node low or high. If this type of scheme is used, you must take into account the various offsets in the output stage. I have found that getting this scheme to initialize a channel's output to exactly zero can be difficult. If you can live with the small errors associated with this type of scheme, it is a low-cost solution.

The last solution is to select an out-

Device	Resolution	Cost
AD5300	8 bits	\$1.25
AD5310	10 bits	\$1.70
AD5320	12 bits	\$2.50

**Table 1**—The AD53xx family offers good performance at rock-bottom prices.





**Figure 4**—This is a sure way to assure a benign power-up state.

put amplifier that has the ability to tristate or shutdown its output. The OPA548 has this capability.

### THAT'S ALL FOLKS

Analog interfaces are found in almost every real-world controller. There are many books dedicated to the subject. Here I have presented a handful of parts and a couple of circuit concepts that have proven useful in past projects. My hope is that they will serve you well in your future designs.

I would like to leave you with one final thought. If you purchase an off-the-shelf PLC or controller, you should endeavor to understand how the I/O is implemented. Purchasing an off-the-shelf controller may reduce your hardware and firmware design cycles. Certainly you will reduce your manufacturing overhead. However, in the end, the reliability of the system is your responsibility.

Obtain as much information about how the PLC company implemented the analog channels. Examine the schematics. Look carefully at the PCB routing. Does the engineering appear sound?

Ask for data on effective resolution. Take the time in your lab to verify the performance the manufacturer advertises.

The PLC vendor's engineers may have made tradeoffs to optimize parameters that may not be the ones your application needs optimized. Try to understand the tradeoffs made by your vendor.

Engineering is a blend of art and science. It's full of tradeoffs and details. Never fear the tradeoffs, and always consider the details.

*Over the last ten years, Bob has designed instrumentation for agronomy,*

*soil physics, and water activity research. He has also designed embedded controllers for a variety of other applications. For more technical resources and articles, visit Bob's online library at [www.engineerbob.com](http://www.engineerbob.com). You may reach Bob with comments and questions at [bob@engineerbob.com](mailto:bob@engineerbob.com).*

### REFERENCES

- [1] R.F. Graf, *The Modern Amplifier Circuit Encyclopedia*, Tab Books, Blue Ridge Summit, PA, 1992, p. 167.
- [2] Analog Devices, *Practical Design Techniques for Sensor Signal Conditioning*, 1999.
- [3] Analog Devices, *1992 Amplifier Applications Guide*, 1992.

### SOURCES

#### AD7887, AD623

Analog Devices  
(617) 329-4700  
Fax: (617) 329-1241  
[www.analog.com](http://www.analog.com)

#### ADS1211, INA118, OPA548

Burr-Brown Corp.  
(520) 746-1111  
Fax: (520) 889-1510  
[www.burr-brown.com](http://www.burr-brown.com)

#### LTC2400

Linear Technology  
(408) 432-1900  
Fax: (408) 434-0507  
[www.linear-tech.com](http://www.linear-tech.com)

#### TLC2543

Texas Instruments Inc.  
(800) 477-8924, x4500  
(972) 995-2011  
Fax: (972) 995-4360  
[www.ti.com](http://www.ti.com)

**Circuit Cellar, the Magazine for Computer Applications.** Reprinted by permission. For subscription information, call (860) 876-2199 or [subscribe@circuitcellar.com](mailto:subscribe@circuitcellar.com).